LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-CONF-458952

Fighting Fire with Fire: Superlattice Cooling of Silicon Hotspots to Reduce Global Cooling Requirements

S. Biswas, M. Tiwari, T. Sherwood, L. Theogarajan, F. T. Chong

October 11, 2010

Graduate Student Workshop 2010
Santa Barbara, CA, United States
October 8, 2010 through October 8, 2010

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Fighting Fire with Fire: Superlattice Cooling of Silicon Hotspots to Reduce Global Cooling Requirements

Susmit Biswas⋆, Mohit Tiwari⋆, Timothy Sherwood⋆, Luke Theogarajan†, Frederic T. Chong⋆
Department of Computer Science⋆,Department of Electrical and Computer Engineering†
{susmit, tiwari, sherwood, chong}@cs.ucsb.edu⋆, ltheogar@ece.ucsb.edu†

I. Introduction

The running costs of data centers are dominated by the need to dissipate heat generated by thousands of server machines. Higher temperatures are undesirable as they lead to premature silicon wear-out; in fact, mean time to failure has been shown to decrease exponentially with temperature (Black's law [2]). Although other server components also generate heat, microprocessors still dominate in most server configurations and are also the most vulnerable to wear-out as the feature sizes shrink. Even as processor complexity and technology scaling have increased the *average* energy density inside a processor to maximally tolerable levels, modern microprocessors make extensive use of hardware structures such as the load-store queue and other CAM-based units, and the *peak* temperatures on chip can be much worse than even the average temperature of the chip. In recent studies, it has been shown that *hot-spots* inside a processor can generate $\sim 800W/cm^2$ heat flux whereas the average heat flux is only $10-50W/cm^2$, and due to this disparity in heat generation, the temperature in hot spots may be up to $30\,^{\circ}\mathrm{C}$ more than average chip temperature.

The key problem processor hot-spots create is that in order to prevent some critical hardware structures from wearing out faster, the air conditioners in a data center have to be provisioned for worst case requirements. Worse yet, air conditioner efficiencies decrease exponentially as the desired ambient temperature decreases relative to the air outside. As a result, the global cooling costs in data centers, which nearly equals the IT equipment power consumption, are directly correlated with the maximum hot spot temperatures of processors, and there is a distinct requirement for a cooling technique to mitigate hot-spots selectively so that the global air conditioners can operate at higher, more efficient, temperatures.

We observe that localized cooling via superlattice microrefrigeration presents exactly this opportunity whereby hot-spots can be cooled selectively and allow global coolers to operate at much more efficient temperatures. Recent advances in processor cooling technologies have demonstrated that thermoelectric coolers (TEC), which use a Peltier effect to form heat pumps, can be used to reduce the temperature of hot spots. By applying a thermoelectric cooler between the heat spreader and the processor die and applying current selectively at the hot spots, heat from the hot-spots can be spread much more efficiently. The ability to implement such thermoelectric coolers on a real silicon device has been demonstrated recently [3], albeit for small prototype chips. The key question then, that needs to be answered before such thermoelectric coolers can be integrated in commodity server processors, is "What is the potential for superlattice microrefrigeration to reduce global cooling costs in data centers?".

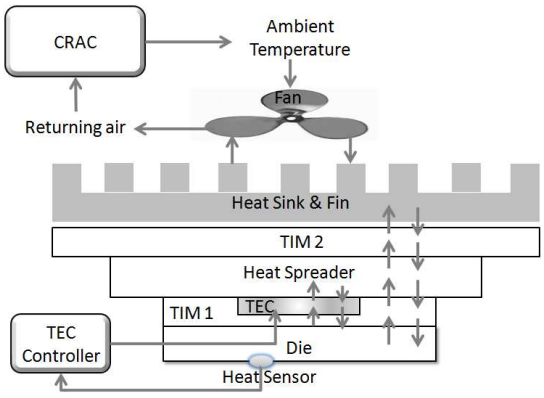

Fig. 1: Heat Flow

In order to answer this question, we present a comprehensive analysis of the impact of thermoelectric coolers on global cooling costs. Our thermal analysis covers all aspects of cooling a server in a data center, and integrates on-chip dynamic and leakage power sources with a detailed heat diffusion

model of a processor (that models the silicon to the thermoelectric cooler to the heat spreader and the heat sink) and finally the computer room air conditioner (CRAC) efficiency, as shown in Figure 1. In Section II, we present the components of the system model.

II. Modeling Chip to Data-center Cooling

In order to quantify the benefits of using TEC based cooling, we need to estimate the peak temperature, which determines the mean time to failure (MTTF), as well as the consumption of power by the microprocessor - taking into account the leakage and temperature feedback loop. We use curve fitting techniques along with ITRS [1] data to model the leakage power consumption in the active layer. We use the finite difference method of computing the steady state temperature by solving heat diffusion equation (1) with boundary condition of convective cooling (equation (2)) using an explicit method which promises greater scalability for fine grain analysis. TEC is modeled as a heat soure at the hot side and a sink at the cold side with pumping rate as shown in equation (3), and we model CRAC efficiency (Coefficient of Performance) with equation (3) as proposed by Moore *et al.* [4] for estimating the cooling power consumption.

$$T^{t+1} = T^t + \left(\frac{k\Delta t}{\rho C_p}\right)\left[\frac{1}{\Delta x^2}\left(M_x + M_y + M_z\right)T^t\right] + \frac{g\Delta t}{\rho C_p} \quad (1)$$

$$M_x(T_{x,y,z}) = T_{x-1,y,z} + T_{x+1,y,z} - 2.T_{x+1,y,z}(M_y, M_z similar)$$

$$T_{surf+\Delta x} = T_{surf} - \Delta x.\frac{h}{k}.(T_{surf} - T_{amb}) \quad (2)$$

$$g_{TEC} = \frac{1}{\Delta x^2}\left(0.5\frac{\alpha^2}{\rho_e}T^2_{cold-side}\right) \quad (3)$$

$$COP(T) = a_2T^2 + a_1T + a_0 \quad (4)$$

$$P_{cooling} = (P_{leak+dyn}/COP_{CRAC,T_{amb}+\Delta}) + Capacity_{TEC} * Area_{TEC}/COP_{TEC} \quad (5)$$

By integrating all the models, cooling power at $(T_{Amb} + \Delta)$ ambient temperature is estimated as equation (4). We perform our experiments by collecting dynamic power density traces of several SPEC CPU2000 benchmarks. Leakage power is modeled using ITRS data [1] and as a 3^{rd} order function of temperature. We developed a tool that solved the heat diffusion equations (1) - (3) using the finite difference method to estimate the total power consumption and the temperature profile of the chip-package. We present the results only for a representative benchmark *apsi* in Section III.

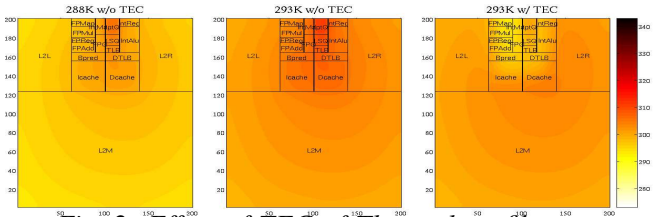

Fig. 2: Effect of TEC of Thermal-profile

III. Results

By increasing the ambient temperature by $5\,^{\circ}\mathrm{C}$ $(288K$ to $293K)$, the peak temperature of the die increases, and by using a TEC layer, the peak temperature is reduced as shown in Figure 2. We experiment with various TEC pumping capacities $(400, 500, 600, 800 W/cm^2)$ and present the result for *apsi* in Figure 3 where bars plotted agains $y1$ axis indicate the cooling power consumption, and MTTF of the hottest point is plotted against the $y2$ axis. Through curve-fitting, we find that TECs allow the CRAC to operate at $6.5\,^{\circ}\mathrm{C}$ higher temperatures $(294.5K)$ which translates into a 25.73% reduction in cooling power, while maintaining the same worst-case on-chip temperatures i.e. without degrading the lifetime of the processor.

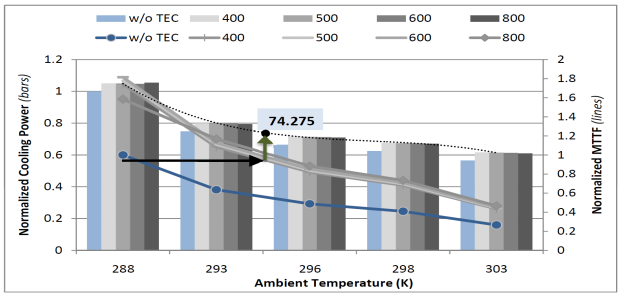

Fig. 3: Effect of TEC on cooling power and MTTF

References

[1] The International Technology Roadmap for Semiconductors. http://www.itrs.net/.

[2] J. Black. ElectromigrationA Brief Survey and Some Recent Results. *IEEE Transactions on Electron Devices*, 16(4):338–347, 1969.

[3] I. Chowdhury, R. Prasher, K. Lofgreen, G. Chrysler, S. Narasimhan, R. Mahajan, D. Koester, R. Alley, and R. Venkatasubramanian. On-chip Cooling by Superlattice-based Thin-film Thermoelectrics. *Nature Nanotechnology*, 2009.

[4] J. Moore, J. Chase, P. Ranganathan, and R. Sharma. Making Scheduling "Cool": Temperature-Aware Workload Placement in Data Centers. In *ATEC '05: Proceedings of the annual conference on USENIX Annual Technical Conference*, pages 61–75, Berkeley, CA, USA, 2005. USENIX Association.

Prepared by LLNL under Contract DE-AC52-07NA27344.